Color Theory and Human Visual System

Optics

- The study of light has 3 sub-fields
  - *Geometric optics*: study of the particle nature of light
  - *Physical optics*: study of the wave nature of light
  - *Quantum optics*: study of the dual wave-particle nature of light and attempt to construct unified theories to support duality, and wave “packets” called photons
- Visualization and graphics most concerned with geometric optics (but need some of the others, too)

Color

- Visible spectrum wavelengths 400-700 nm
- A given color has some distribution of these wavelengths
- Intensity of each wavelength determines contribution to color

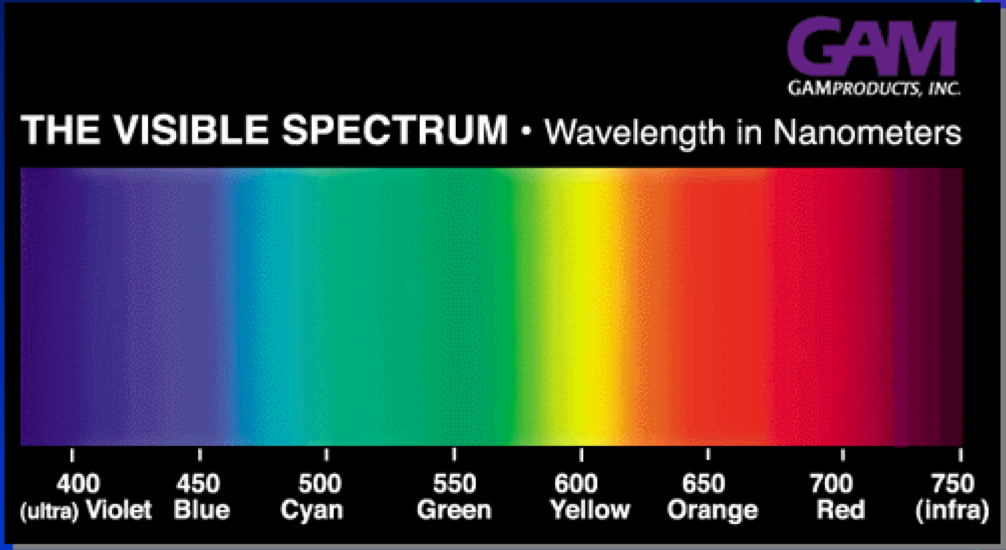

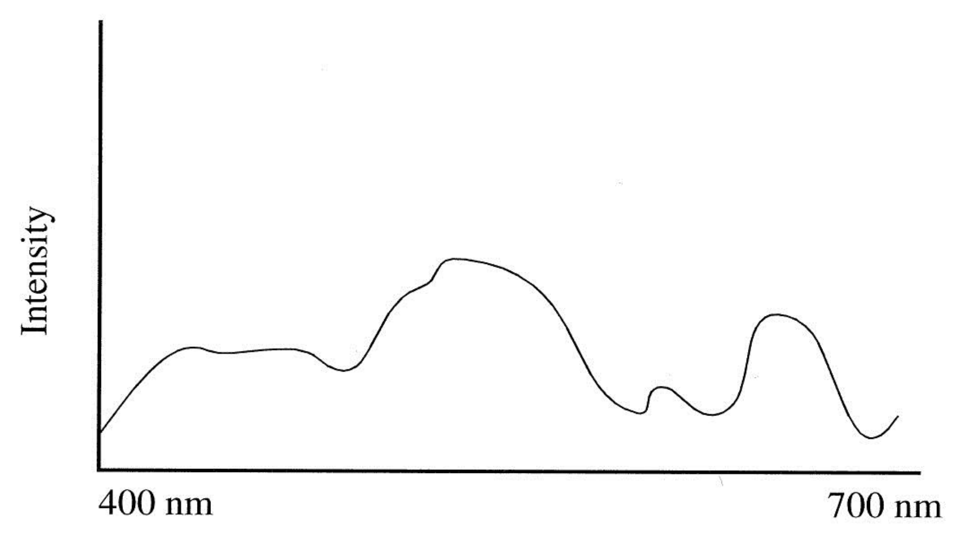

Color

- Color is determined by the wavelength of visible light
- Still use geometric optics
- But need to account for wavelength in reflectance (BRDF) and index of refraction
- What natural phenomena can you think of that are wavelength dependent?

Wavelength Sampling

- We could try to compute image for every possible wavelength and then combine
  - Would take forever
- Sample a representative set of wavelengths
  - How many samples?
  - Where?

Where to Sample?

- Photometry tells us that some wavelengths are more important than others to human perception.
  - Human response curve looks something like this:

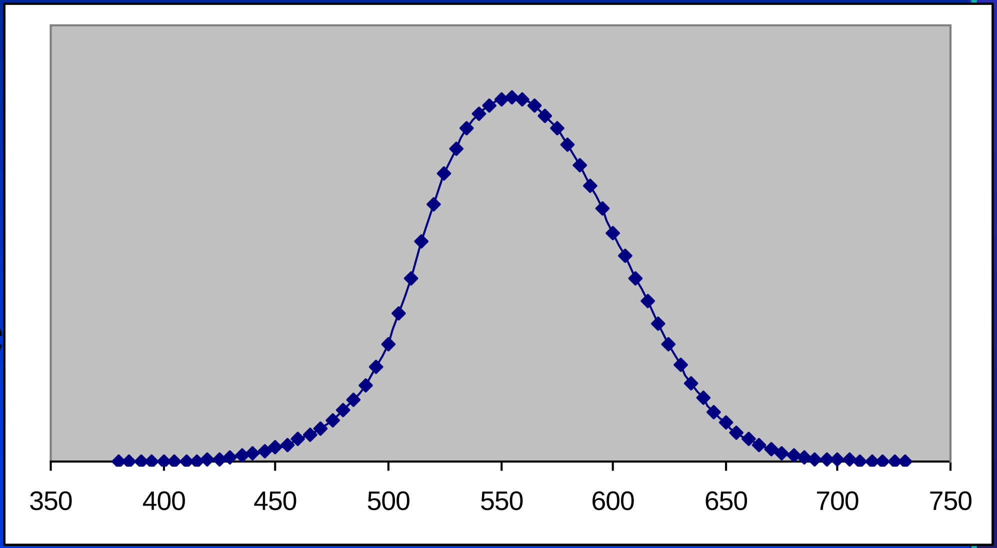

Where to Sample?

- So, pick a few samples wavelengths.
  - Compute an image for each.
  - Reconstruct with basis functions.
  - Weight of each sample determined by human response curve.
  - (Also need color space transformations)

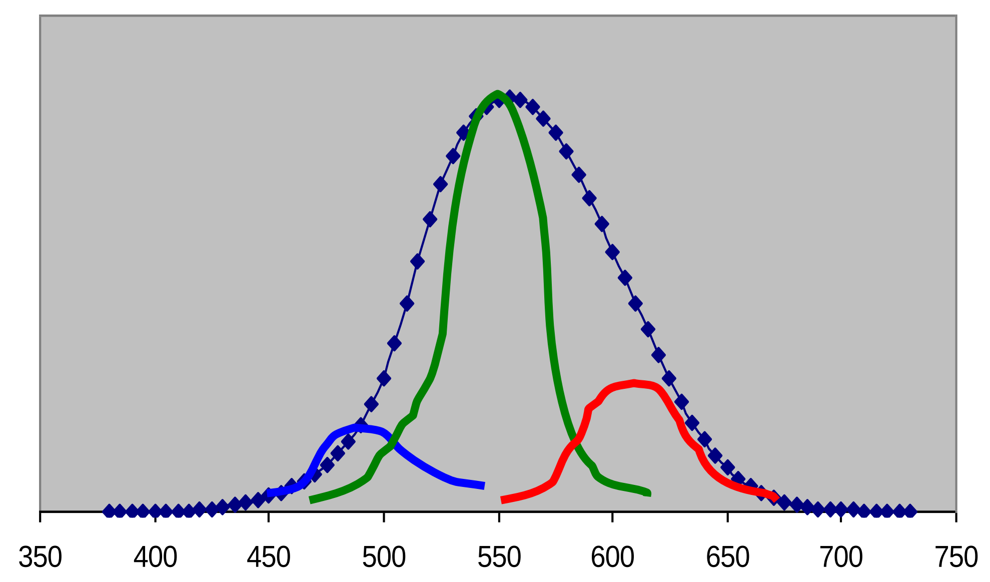

Human Visual System and Color Theory

- Today: human visual system
- Human eye
- Color models
- Color perception

Human Perception (Visual) System

How do we perceive the visible world?

1. Light enters eye and eventually strikes lens
2. Muscles expand and contract to focus light on retina at the back of the eye

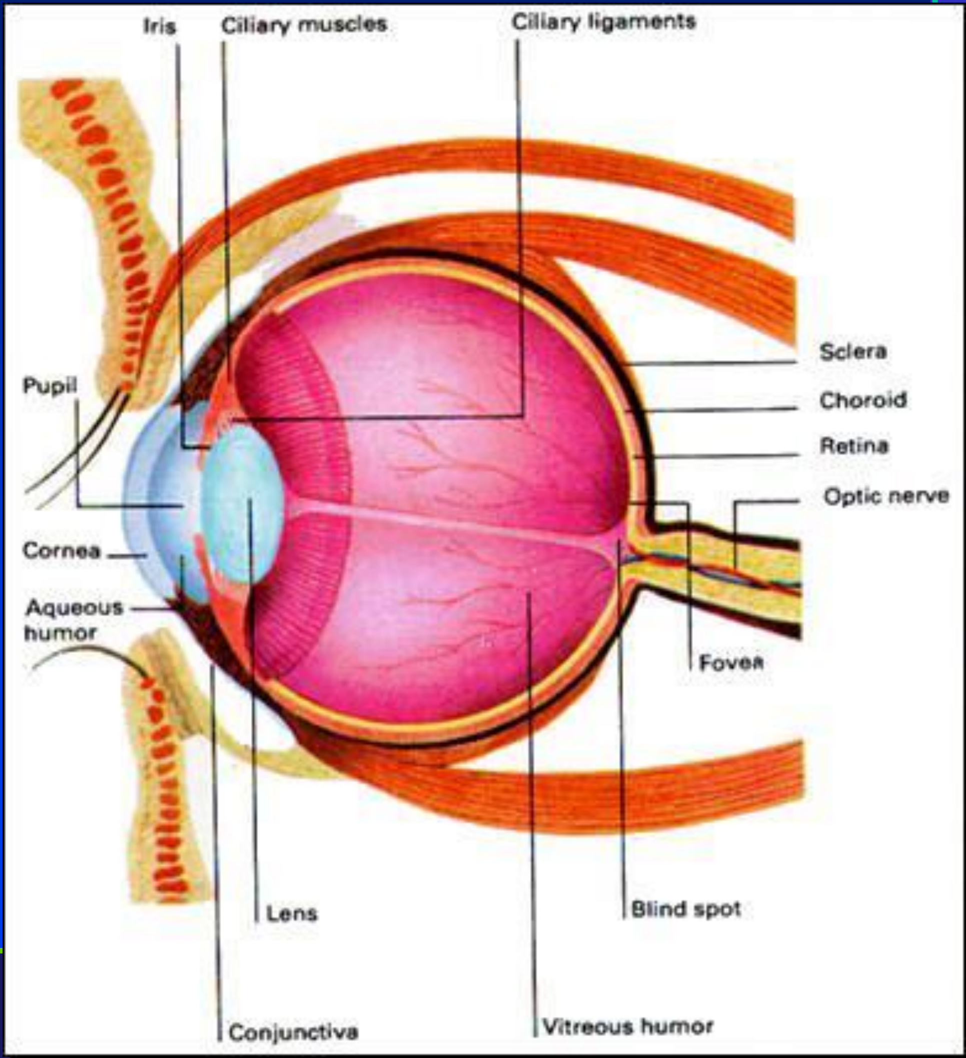

Human Perception (Visual) System

3. Retina senses light and contains **cone cells** and **rod cells**
4. Retinal nerve fibers connect to optic nerve, which carries signals to brain, where they are interpreted as an image

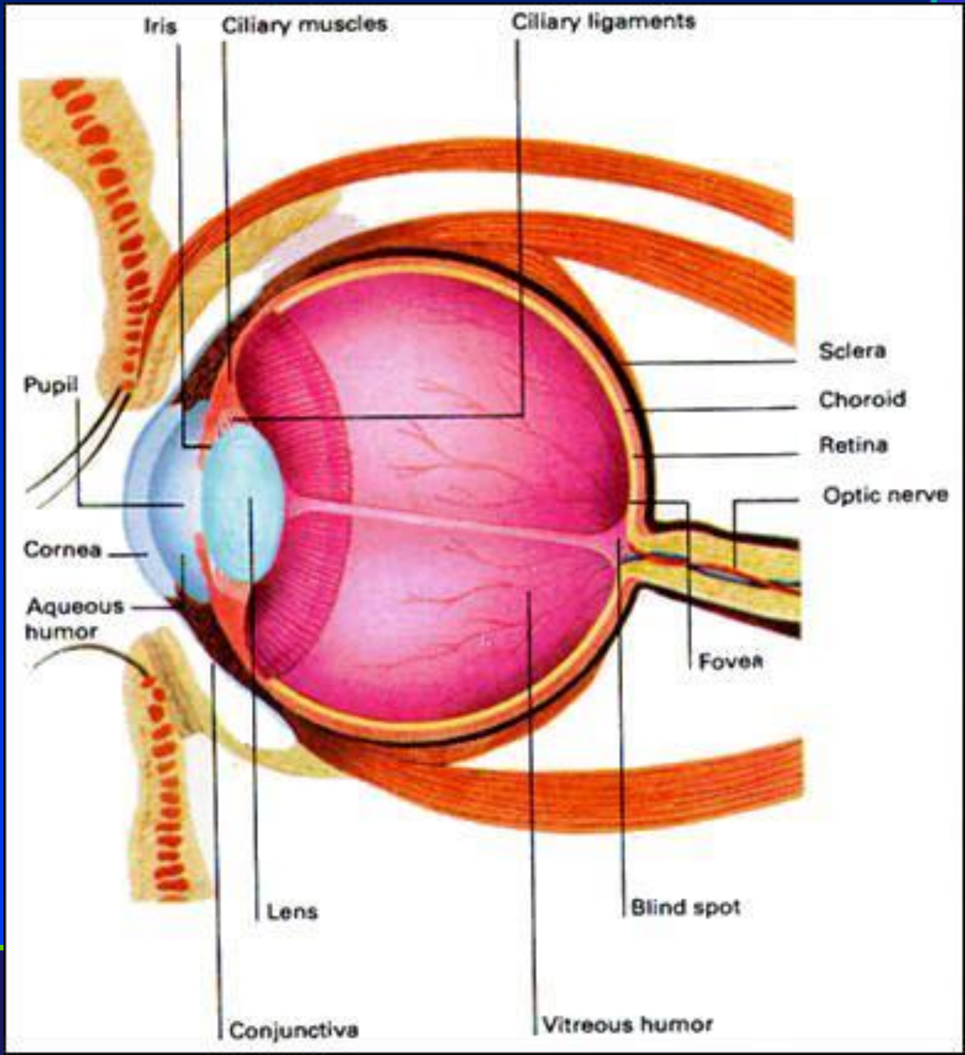

Rods

- Spread all over retina surface
- 75-150 million total
- Low resolution
- Do not detect color
- But very sensitive to low light
- This is called **scotopic vision**

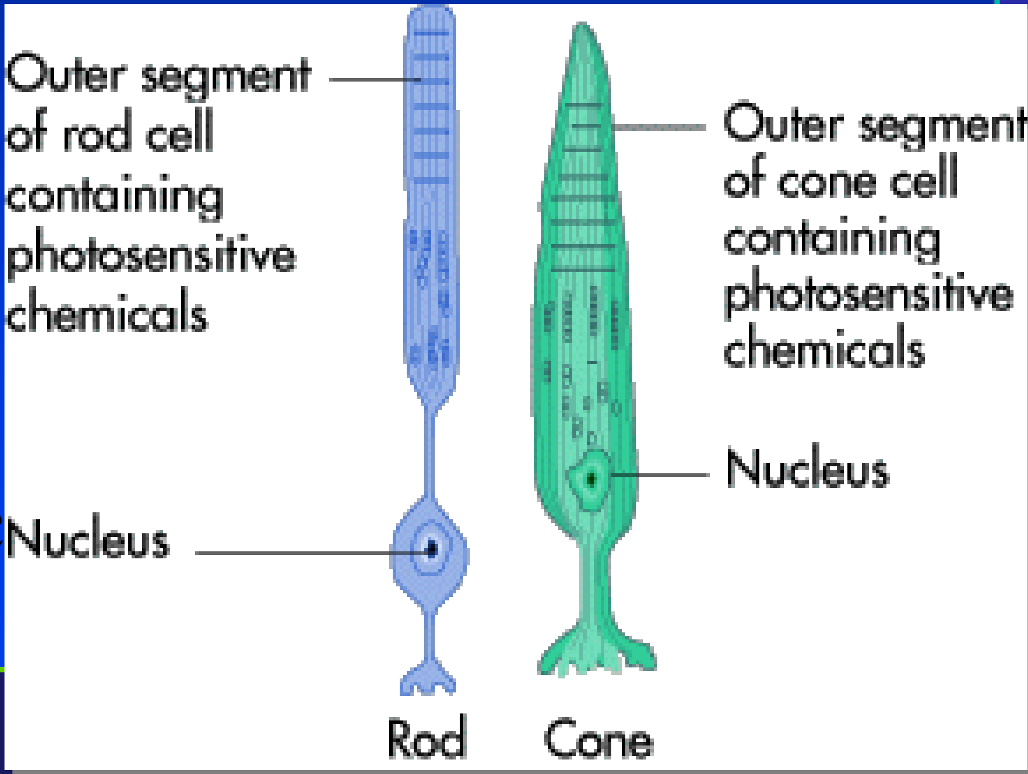

Cones

- A dense array of cells around central portion of retina
- 6-7 million total
- High-resolution, detect color
- Require bright color (This is called **photopic vision)**

- Our eyes, therefore, have a fixed resolution, much like a computer screen
- Our brains permits us to perceive the images as continuous even though they really aren't
- This happens through a process known as *interpolation*

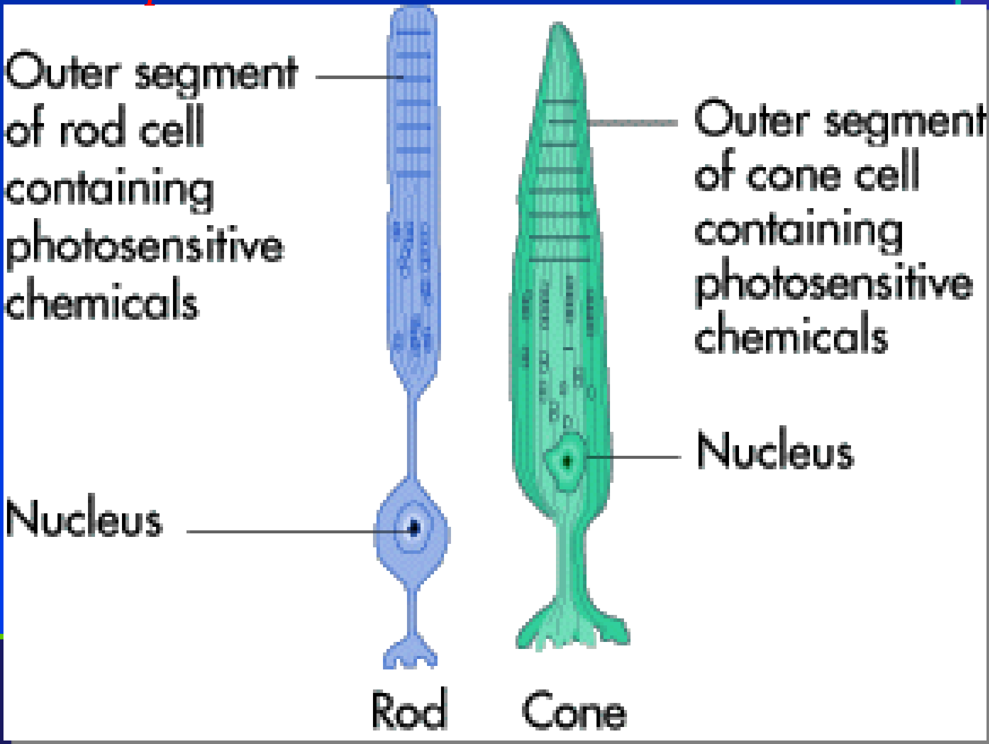

Cones

- Detect color on the retinal surface
- Can be divided into “red” cones (~60%), “green” cones (~30%) and “blue” cones (~10%)
- A complex mixing process takes place inside the brain to generate the final color
- Graph of receptor sensitivities of each type of cone:
- Question: what kinds of colors are we best able to distinguish from each other?

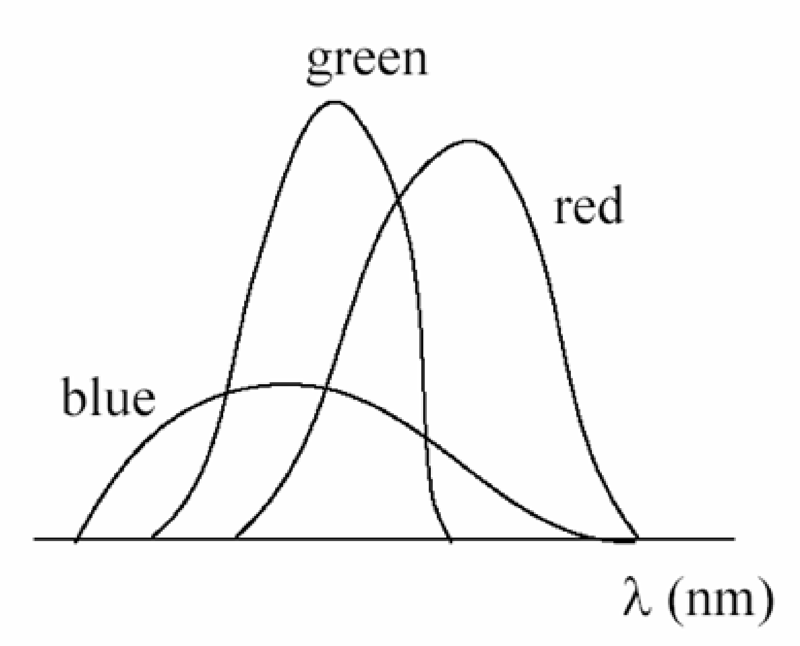

From Human Perception to Digital Images

- Now that we know how the eye perceives light intensity and color, let’s take a “look” at digital images
- We’ll see what is required in image generation and display in order for the eye to perceive intensities and colors correctly/effectively
- Based on what we learn, we will derive some general principles of color-enabled visualization that will help you in your own visualization-related efforts
  - when is color helpful for understanding
  - when it is necessary for understanding
  - when it distracts from or hampers understanding

Pixels and Images

- Image = 2D matrix or grid of pixels
- Pixel = a dot of light (possibly colored)
- Image resolution = number of pixels along each matrix

res = 300^2 pixels

res = 150^2 pixels

res = 75^2 pixels

res =37^2 pixels

- What is the resolution of a television?
- Why doesn’t it look blocky?

Pixels and Images

- Each pixel has a value:
  - a single value if gray-level image
  - a triple RGB (red, green, blue) if color image
  - we call these the color *channels*
  - sometimes a pixel has four channels: RGBα or RGBA, where the A stands for opacity (opposite of translucency)
- Each pixel is represented by a number of bits:
  - typical is 1 byte = 8 bits, which gives $2^8 = 256$ gray levels
  - a color pixel with 2^8 levels per color channel needs $3 \times 8 = 24$ bits of storage
  - 24-bit color is sometimes called "true color" in graphics-intensive applications because it can represent so many colors and the human eye cannot perceive anything missing

How Many Bits Do We Need?

- Top left: original image (8-bit)
- Top right: lower 4 bits dropped
- Lower left: (top left - top right)
- Lower right: enhanced version of lower left

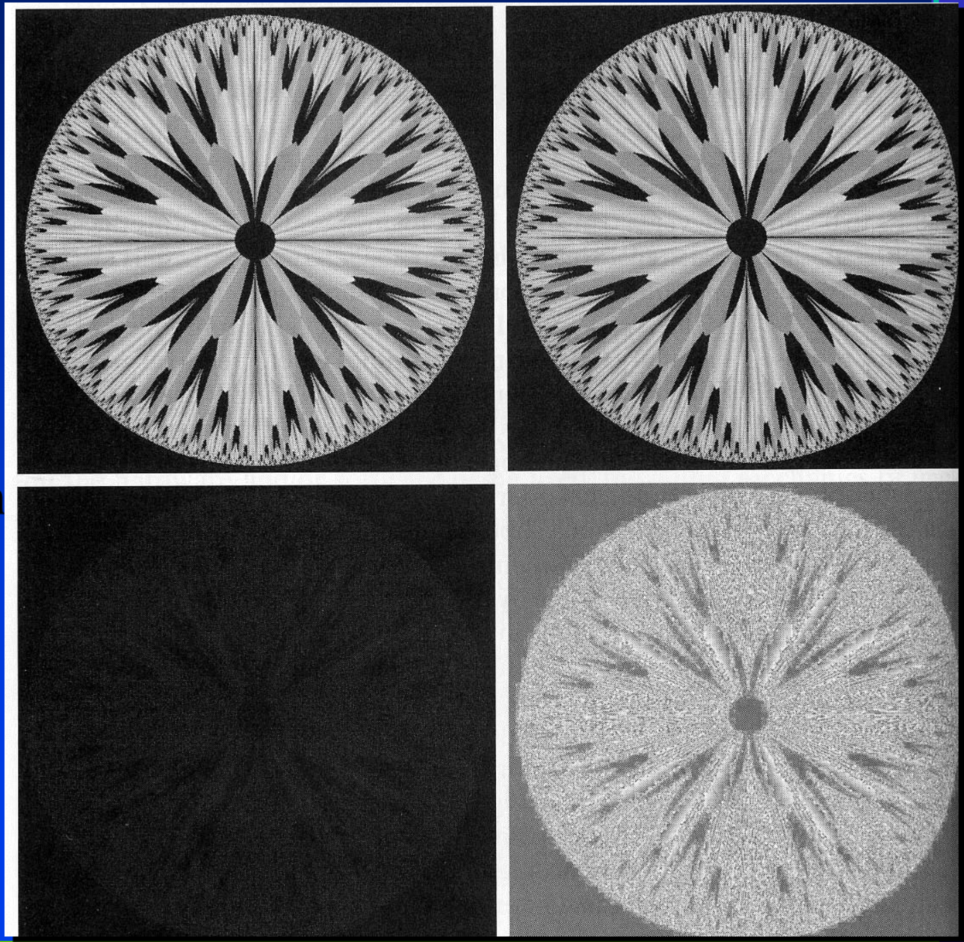

- Even in this day of high speed networks and huge hard disks, data reduction is important, especially with increased use of CT, MRI and other volumetric scanning modalities
- e.g., Visible Human: 13 GB, Visible Woman: 40 GB
- Now multiply this by 300 million residents of U.S.

Display Technology – Monochrome (Gray-scale ONLY)

- A monochrome display monitor has one electron gun that fires a beam of electrons to the inside of the screen
- This electron beam scans the screen in horizontal scanlines, top to bottom, at rates of 30-60Hz or more
- When we wish to draw an object on the screen (e.g., an animated character), the geometric description (e.g., triangles) has to be converted into pixels
- This process is called *scan conversion*
- The pixels to draw on the screen are stored in a special memory area called the *framebuffer*

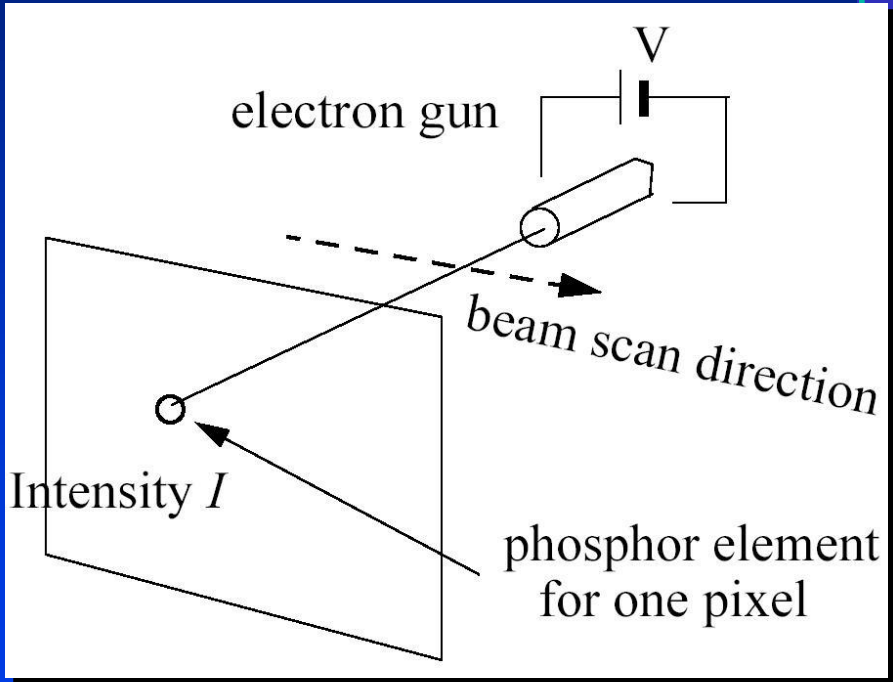

Display Technology – Monochrome (Black & White ONLY)

- The electron beam is steered to the pixels by two orthogonal pairs of charge plates
- The plates bend the beam (with electromagnetism) so that it hits a particular position on the screen
- Each screen pixel consists of a phosphor, which glows at some intensity *I* when hit by a beam
- The phosphor glows for a certain period of time known as the *persistence* of the phosphor (10-60 microseconds)
- Phosphor intensity *I* is a function of the strength of the electron beam (number of electrons, voltage *V*)

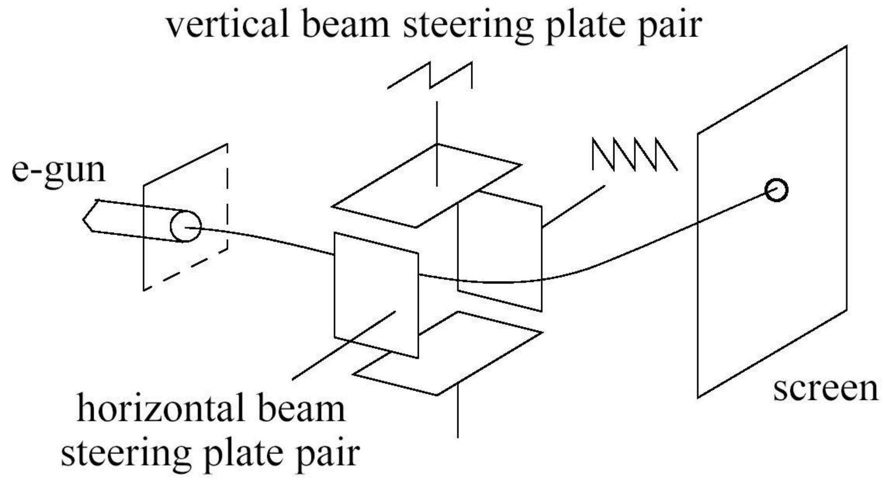

Display Technology – Monochrome (Black & White ONLY)

- The amount of energy that the electron gun delivers to each phosphor depends on the value of the image pixel that is to be displayed there
- For instance, if the pixel should be black, then beam is turned off while it passes over that position on the screen
- Question: if phosphors glow only 10-60 microseconds, but we refresh the screen only once every 1/60th of a second (~17000 msec), how come we don't see obvious flickering?
- Answer: *persistence of vision*. Your eye has something akin to a "memory". When you see something, the image actually persists in your sight for a brief moment even after the image disappears. In other words, the human eye isn't "fast enough" to notice the flicker.
- This phenomenon is not well understood and is the cause of some controversy among those who study this sort of thing

A Closer Look: The Anatomy of an Image

- Do not think of pixels being solid squares
- Rather, each pixel is an infinitesimal small dot, a sample point of zero size
- So why don't we see these images as small disconnects dots on monitors and printed paper?
- Interpolation

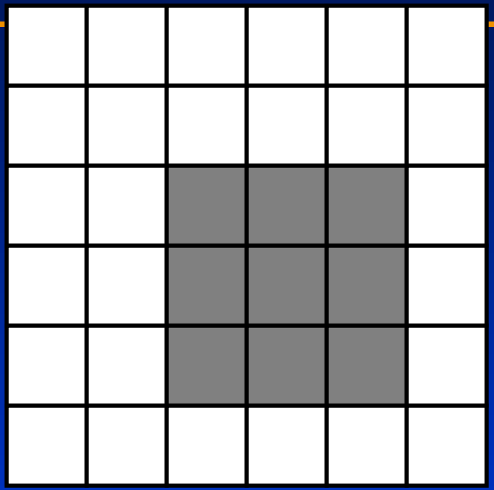

NO

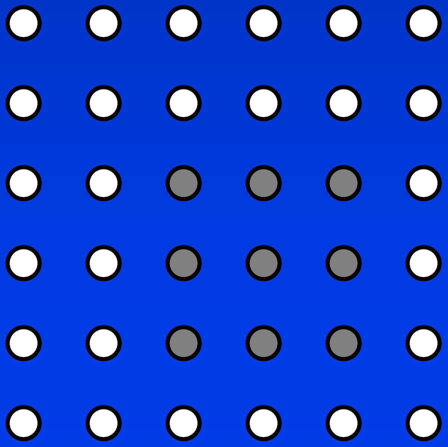

YES

A Closer Look: The Anatomy of an Image

- A monitor or a printer “splats” a pixel to the screen or onto the paper, respectively
- Each pixel assumes the shape of a smooth Gaussian function
- The Gaussians of neighboring pixels overlap slightly and blend together
- This makes the image appear smooth and continuous on a display or paper
- But always be aware that the image is only stored as a number of small colored dots at discrete positions on a square grid

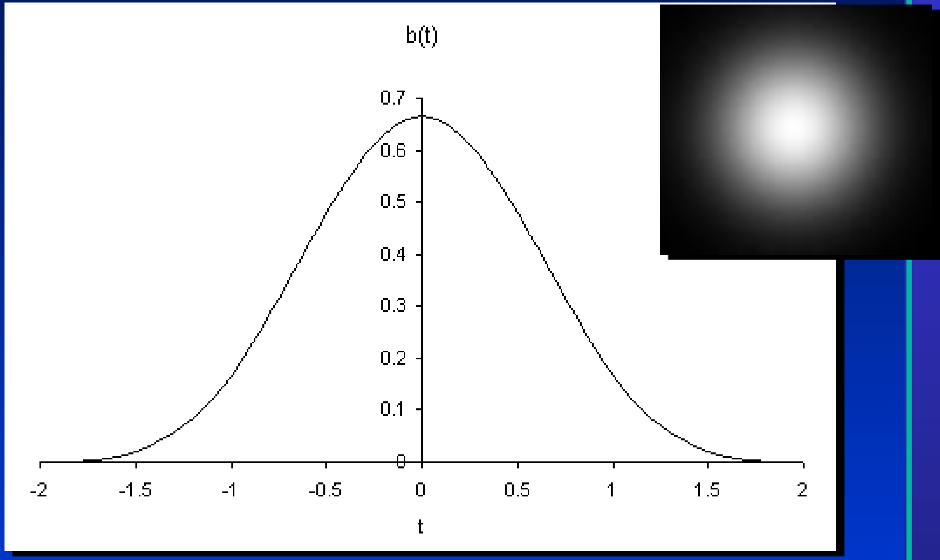

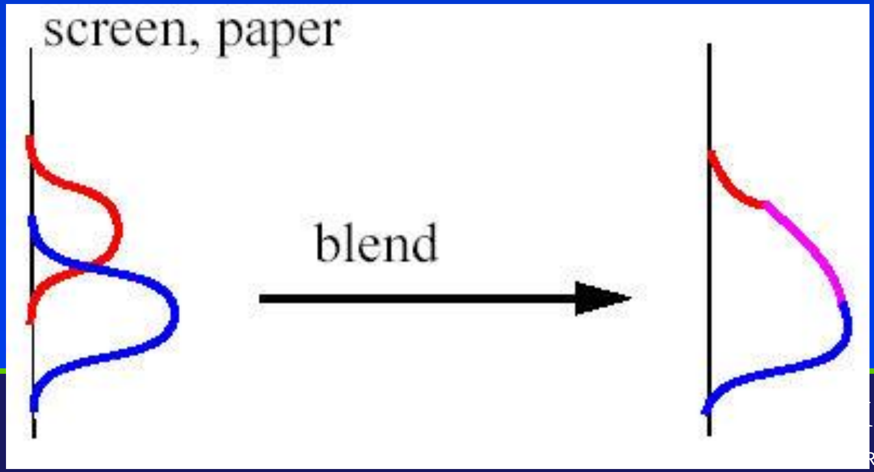

Perception of Light Intensity

- The perceived intensity difference of two intensities is a function of the ratio of the two intensities
- Consider a 50-100-150 Watt light bulb: a 50→100 transition is more noticeable than 100→150
- Suppose our monitor has (assume a monochrome B/W monitor) n+1 intensity pixel graylevels, a minimum intensity I_0, and a maximum intensity $I_n = 1.0$
- We would like steps in brightness at all pixel levels *j*: $\frac{I_{j+1}}{I_j} = r = const$
- Thus,

$$I_1 = rI_0 \quad I_2 = rI_1 \quad \cdots \quad (I_j = r^j I_0) \quad \cdots \quad I_N = r^n I_0 = 1.0$$

Perception of Light Intensity

$$I_1 = rI_0 \quad I_2 = rI_1 \quad \cdots \quad (I_j = r^j I_0) \quad \cdots \quad I_N = r^n I_0 = 1.0$$

We can calculate

$$r = \left(\frac{1.0}{I_0}\right)^{\frac{1}{n}}$$

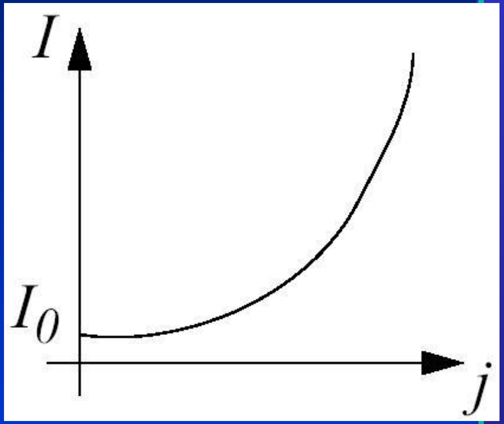

- Experiments reveal that a digital reproduction of an image appears continuous when $r \leq 1.01$
- Thus, the minimum number of intensity levels is:

$$n_{\min} = \log_{1.01}\left(\frac{1.0}{I_0}\right)$$

How Many Intensity Levels Do We Need?

- Given this n_{min}, the eye will just be able to distinguish between I_j and I_{j+1}
- Greater n_{min} are unnecessary since the eye will not be able to distinguish two consecutive gray levels
- Display monitors need 400-530 levels (but usually have only 256)
- Photographic slides have a lower I_0 and thus need more levels (about 700)
- Newspapers have large I_0 and thus can get away with fewer levels (about 230).
- Why do they have large I_0?

8 bit-planes = 256 levels

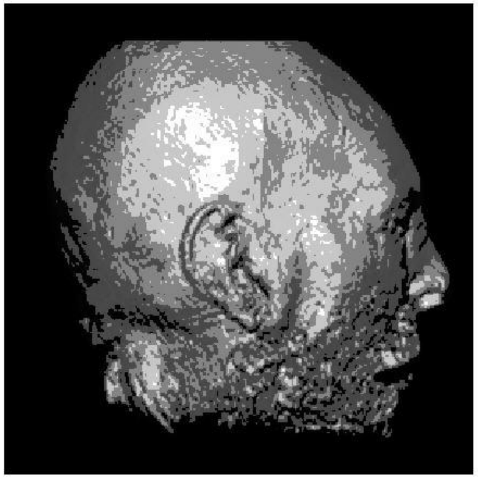

3 bit-planes = 8 levels

Gamma Correction

- We'll see now how all this applies to computer displays
- The monitor brightness *I* is related to the number of electrons, *N*, that strike the phosphors: $I = kN^{\gamma}$
- The number of *photons* is non-linearly related to the voltage *V* at the electron gun: $I = KV^{\gamma}$

$$V = \left(\frac{I}{K}\right)^{\frac{1}{\gamma}}$$

- We need to correct for this non-linearity:
- This is called *gamma correction* (for standard monitors, γ = 2.2 – 2.5, for Trinitron CRTs γ = 1.4 – 1.8)
- Let's look at a practical example to illustrate how this process actually works

Gamma Correction

- Suppose our pixels have an intensity range of *j* = [0, 255]
- We would like to know the voltage *V* that we have to apply at the electron gun
  1. Calculate the intensity that yields uniform perception:
  2. From that, calculate the voltage with proper gamma correction:

$$I = r^j I_0$$

$$V = \left(\frac{I}{K}\right)^{\frac{1}{\gamma}}$$

- Most monitors have gamma correction hardware that performs these steps automatically, but often allow changes to be made manually via software
- If done automatically, the hardware stores the pixel values *j* in the framebuffer and then translates them into the proper voltage

Color Receptors

- Cones detect color on the retinal surface
- *Tristimulus theory*: “Red” cones (~60%), “green” cones (~30%) and “blue” cones (~10%)
- Mixing process takes place inside brain
- Graph of visible spectrum each type of cone is sensitive to:

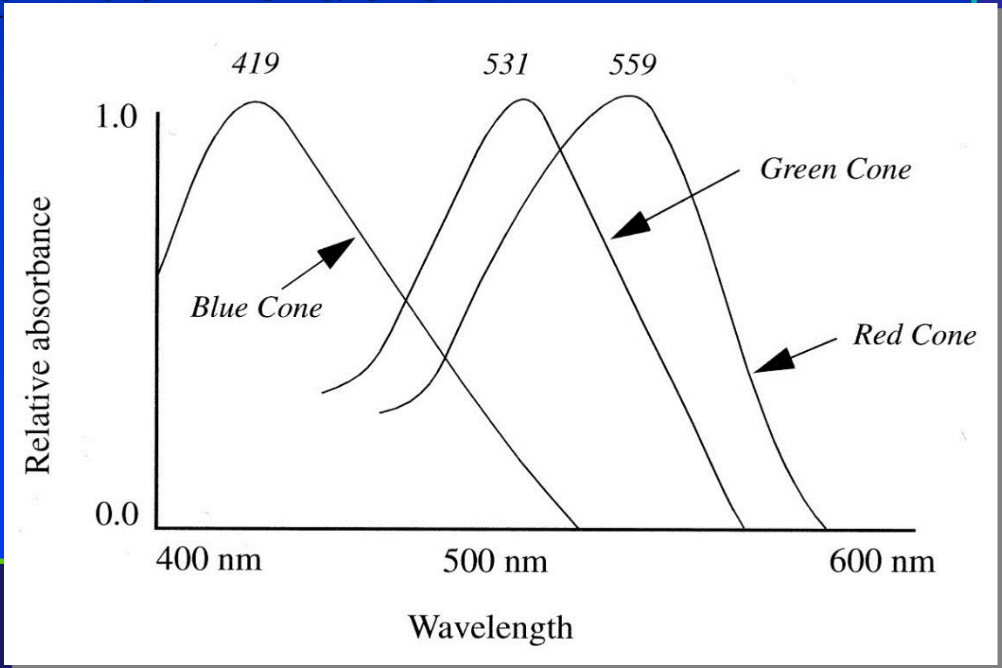

Color Primaries

- Now let's expand our discussion of image generation to color
- As we saw earlier, it is believed that the eye has three kinds of color receptors, one each for red, green and blue
- This is known formally as the *tristimulus theory*
- This means that we can generate nearly any color experience by mixing these three primary colors together on the retina (cones)
- The figure at the right shows the R, G, and B primaries needed to generate a color:
- Notice anything a bit odd about these graphs?

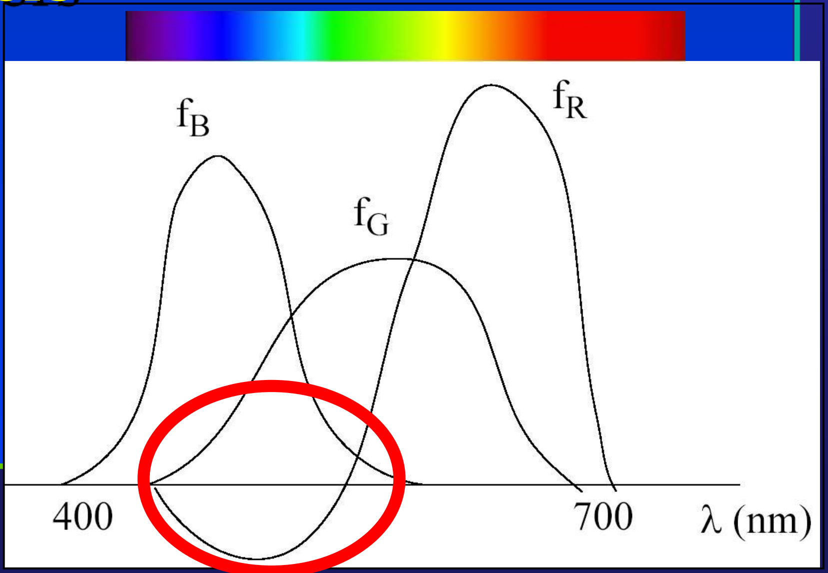

Color Primaries

- The negative part of the graph indicates that the brain does some sort of negative mixing of colors
- This cannot be done with a display monitor that has R, G and B electron guns (note: each gun still fires electrons, not colored light)
- Around 500 nm there are certain colors that cannot be displayed on a computer monitor because a negative amount of red would be required
- Let’s look at how color CRTs work to generate color images

Display Technology - Color

- A pixel in a color image has three components (channels): red, green, and blue
- A color monitor has three electron guns: one for red, one for green, and one for blue
- The beams scan the screen in horizontal scanlines, as before
- But now, each screen pixel consists of a phosphor triple: one glowing red, one green, and one blue
- The glowing phosphor triples blend together to form the color that is encoded in the RGB triple
- The amount of energy that the electron guns deliver to each phosphor depends on the RGB value of the image pixel that is displayed there

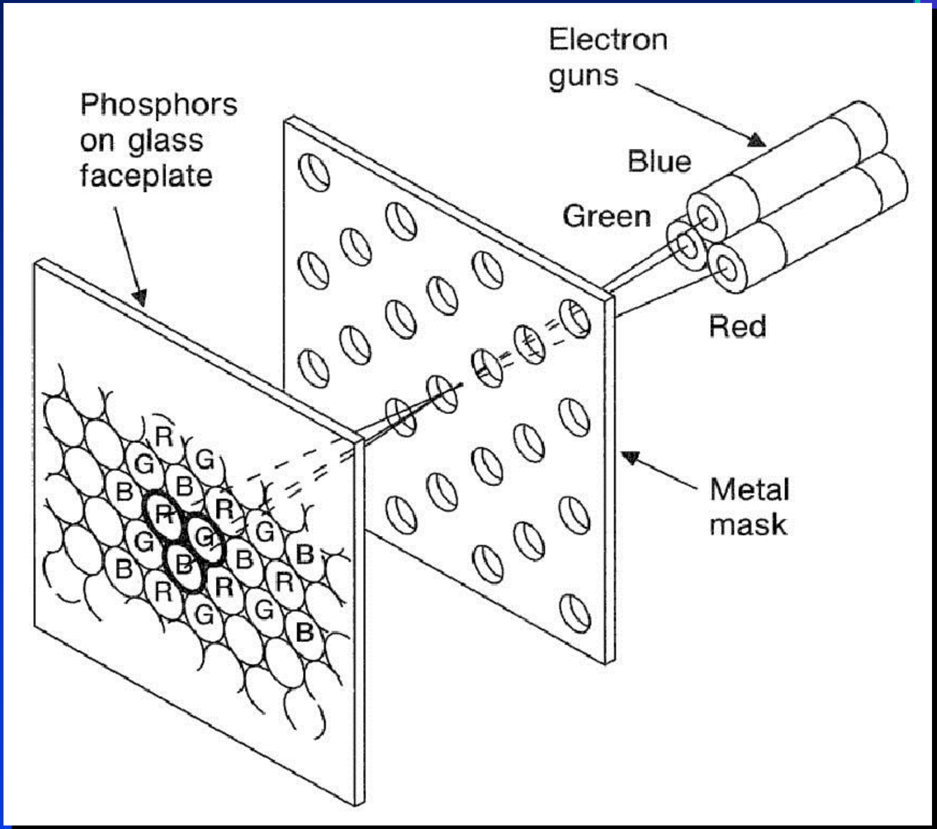

Representation of Color – RGB Color Space

- So how are colors represented in a computer?
- One convenient choice is the RGB *color space*
- R, G, and B in the range [0.0, 1.0]
- black = (0,0,0), white = (1,1,1)
- Complementary colors lie opposite each other on the cube
- Although simple and very fast to use in hardware, the RGB color model is unintuitive when you want to change the saturation (how much gray) and brightness (how much white) is present in the image

- The RGB color cube:

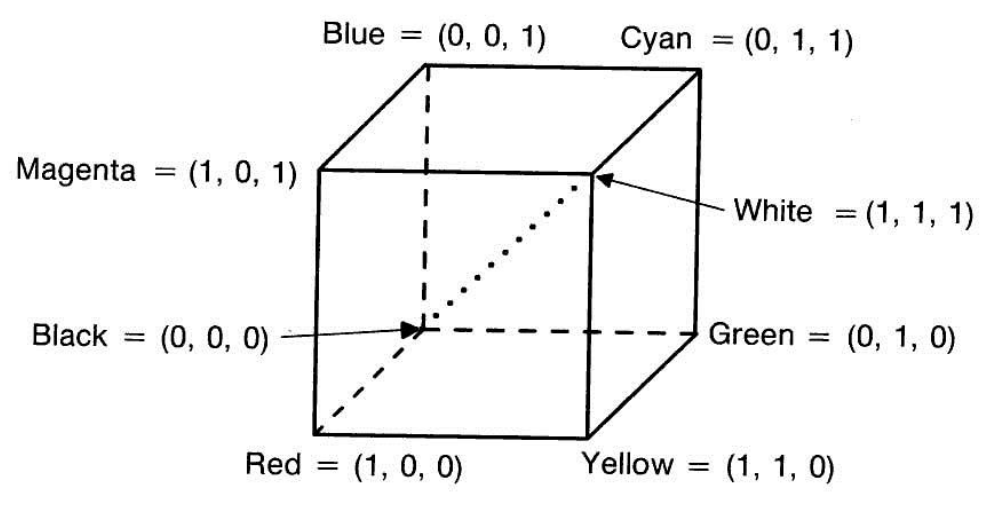

Computer Representation of Color

- Each screen pixel is combination of R, G and B light
- Three *color components* determine perceived color
- *RGB color model*
- R, G, and B in range [0.0, 1.0]

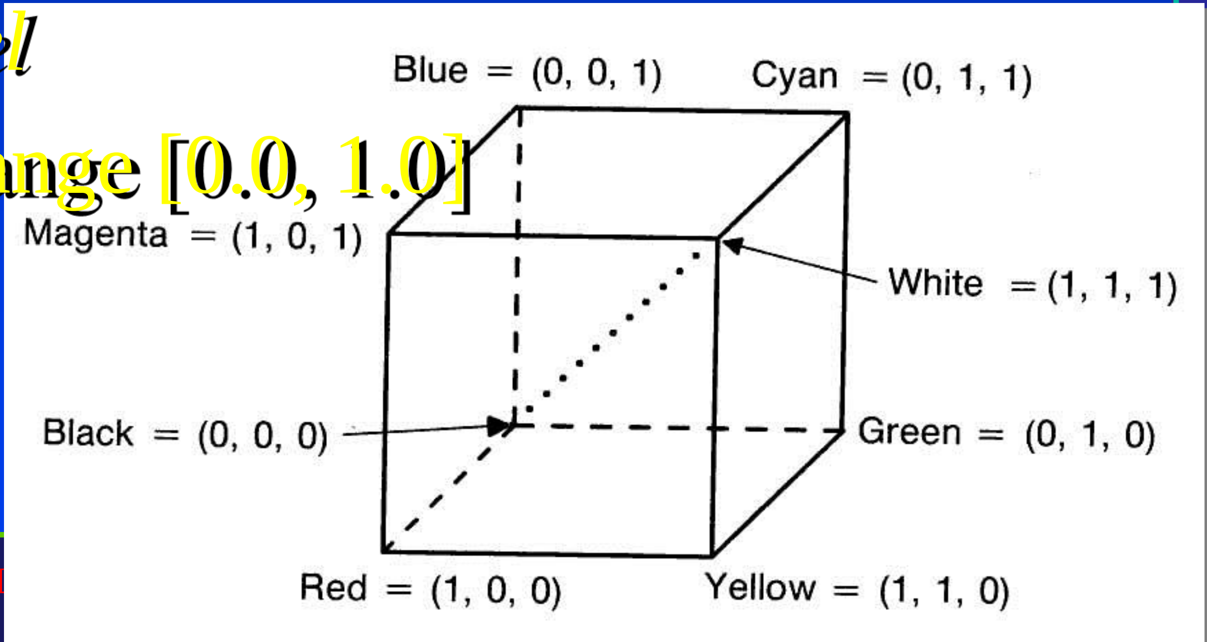

RGB Color Arithmetic

- R + G = ?
- B + G = ?
- R + Y = ?

- R + C = ?
- G + M = ?
- B + Y = ?
- B + W = ?

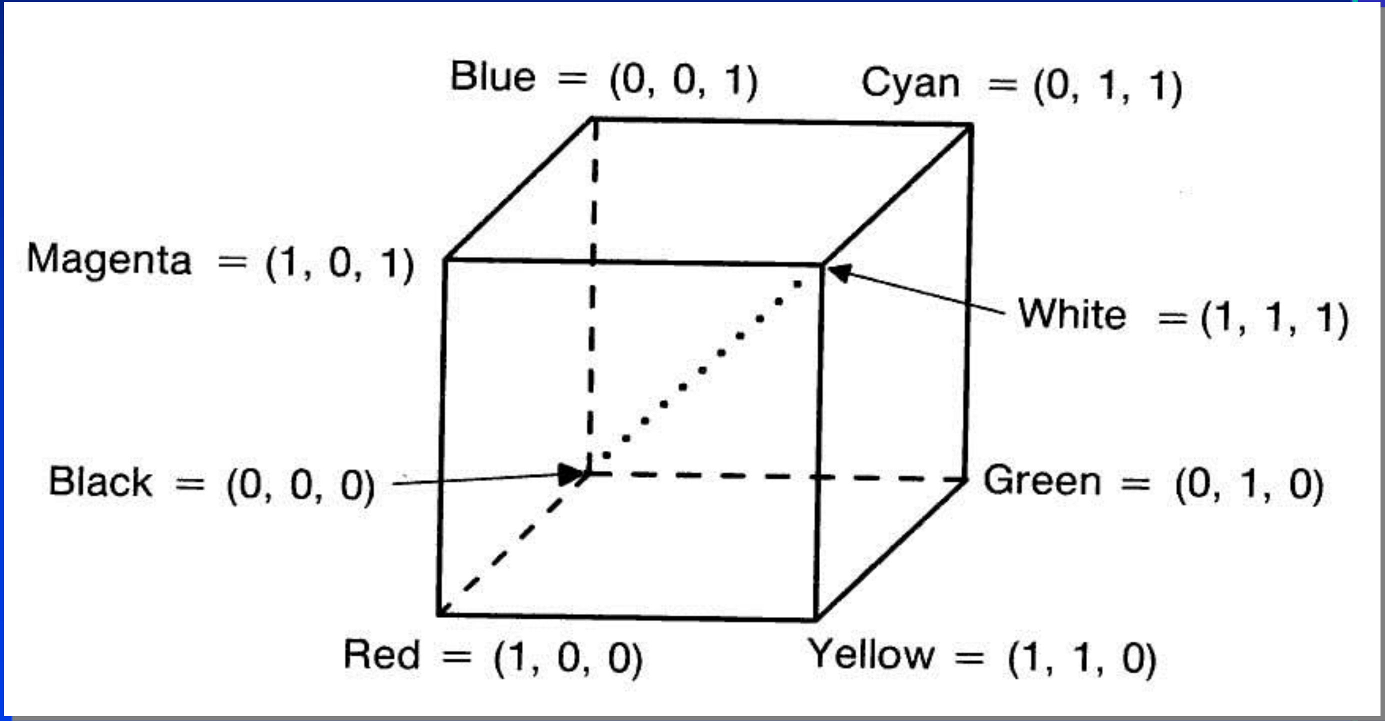

RGB Color Model

- Good: simple, easy hardware implementation
- Is it intuitive?
- How would you make a washed-out green?
- How would you make a bright blue?

Representation of Color – CMY Color Space

- The CMY color space is the complement of the RGB color space and is used primarily for printing (especially if you add black: CMYK)
- The conversion between RGB and CMY is linear:

$$\begin{bmatrix} C \\ M \\ Y \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

- With CMY, the more C, M and Y you add, the darker the color gets
- With RGB, it's just the opposite
- RGB: *additive mixing with light*
- CMY: *subtractive mixing with pigments*

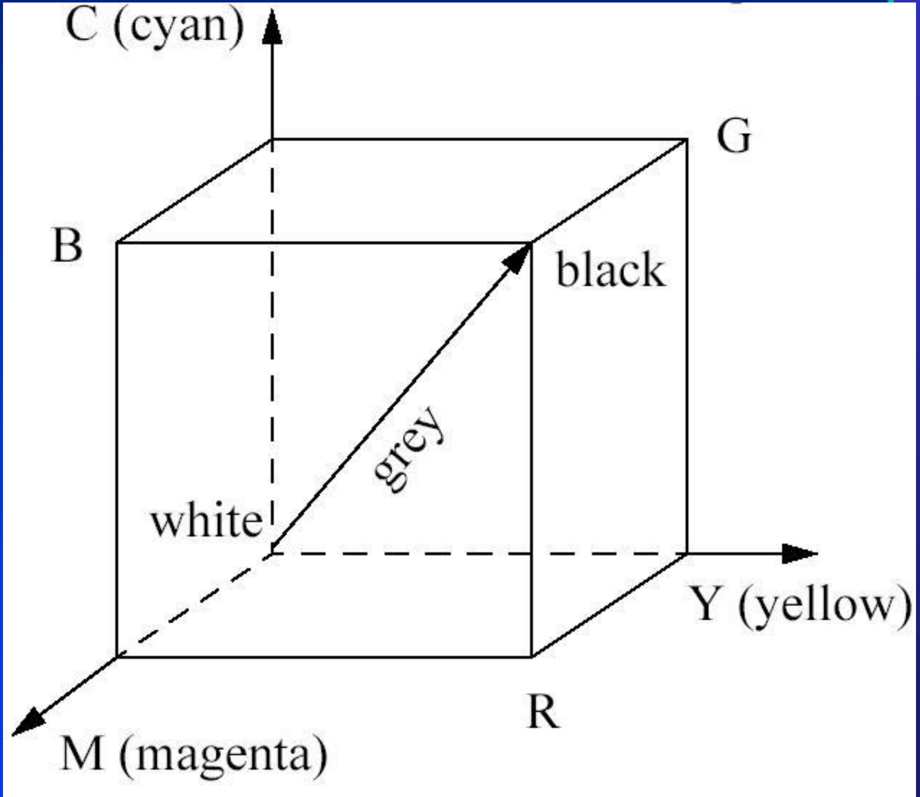

Representation of Color – HSV Hex Cone

- A more convenient color space for specifying colors is HSV
- Hue, saturation, value (brightness)
- S, V in range [0.0, 1.0]
- H in range [0^o, 360^o]
- Value measures the brightness, or white content
- Saturation measures how vivid or washed out the color is
- Changes in S and V are linear with respect with the HSV color model, but non-linear in RGB
- Easier to specify color in HSV and then translate it into RGB for the hardware

- See Foley pp. 592-593 for algorithms to convert between RGB and HSV

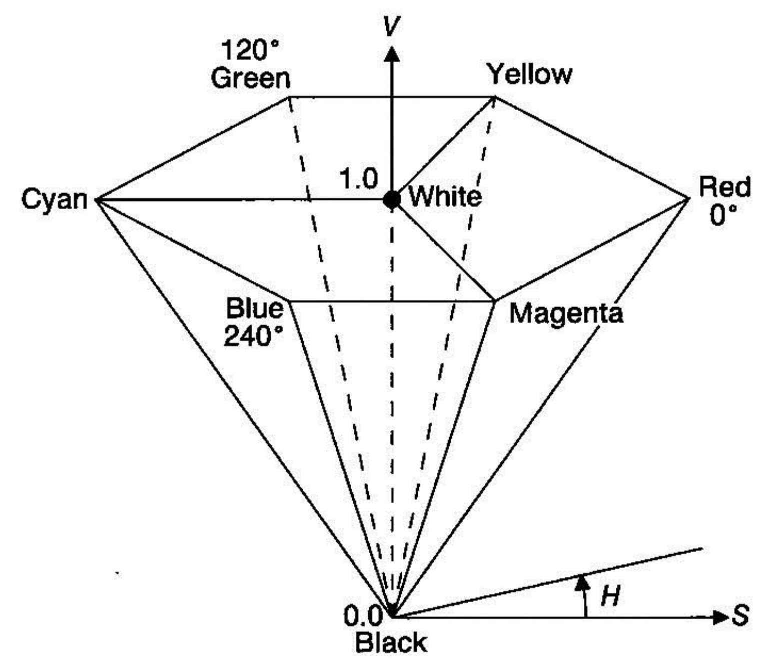

HSV Color Model

- RGB: good for hardware, bad for human use
- Hard to change saturation and brightness
- *HSV color model* – more intuitive
- H = hue
- S = saturation
- V = value (brightness)

HSV Color Model

- *Saturation* measures vividness of color
- Distance from central axis
- *Value* measures brightness
- S, V in range [0.0, 1.0]
- H in range [0°, 360°]

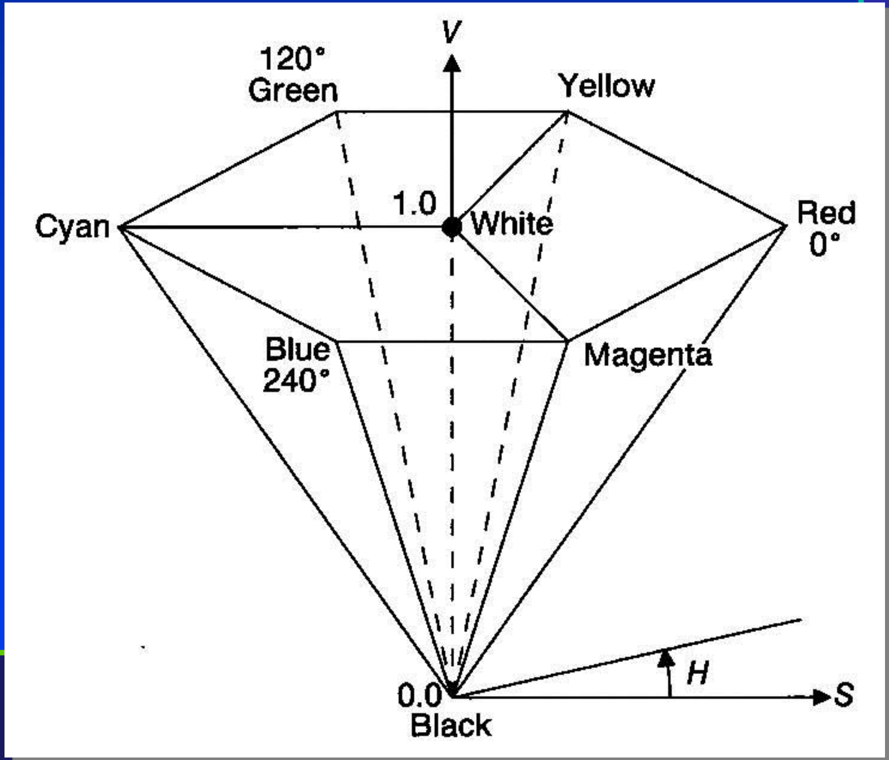

HSV Color Model

- HSV very intuitive
- To brighten a color, increase V
- To wash it out (make grayer), decrease S
- How do we change H?
- Change angle
- Sequence of colors around cone:

$$R \rightarrow Y \rightarrow G \rightarrow C \rightarrow B \rightarrow M \rightarrow R$$

HSV Color Model

- The *hue* is the color
- Specified as angle around the *HSV (hex) cone*

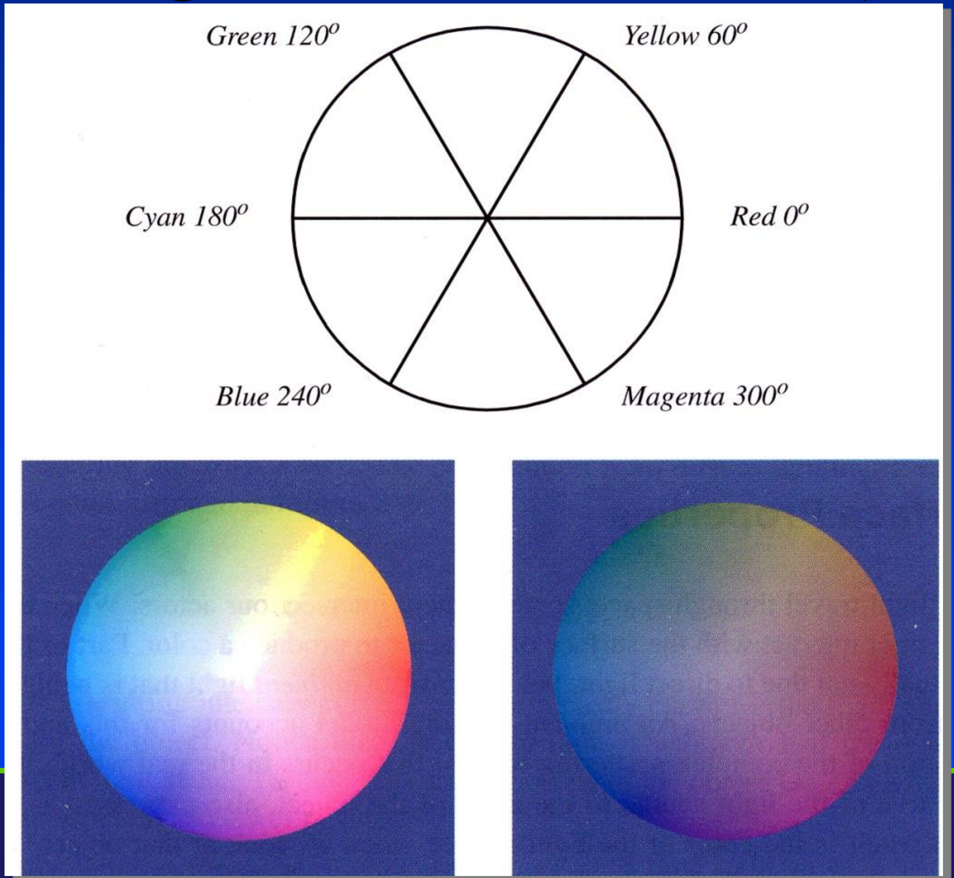

HSV Examples

- Which component of HSV are we increasing in the left image? Right image?

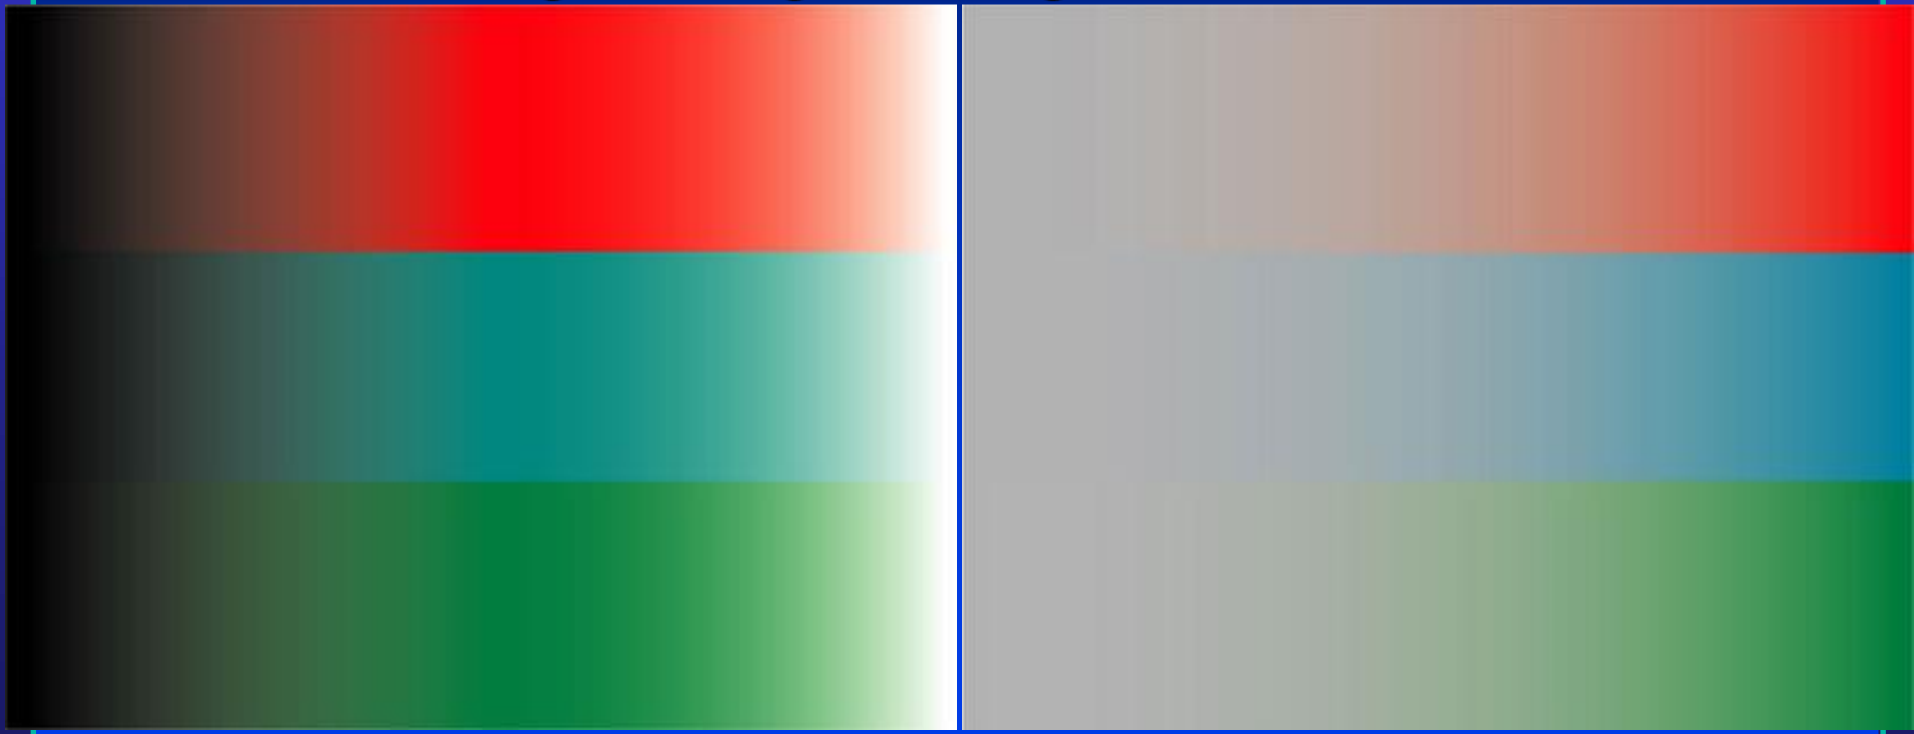

HSV Examples

- Lightness corresponds to how much light appears to be reflected from a surface
- Basically, it measures how much white or black is present
- Saturation is the degree of color intensity associated with a color’s perceptual difference from a white, black or gray of equal lightness

RGB vs. HSV

- Can use HSV in software, but convert to RGB for display
- Easy to convert between RGB and HSV
- Code on Web

Color	**RGB**	**HSV**
Black		
White		
Red		
Green		
Blue		
Yellow		
Cyan		
Magenta		
Sky Blue		

RGB vs HSV

- Can use HSV in software, but convert to RGB for display
- Easy to convert between RGB and HSV
- Code on Web

Color	**RGB**	**HSV**
Black	0,0,0	*°,*,0
White	1,1,1	*°,0,1
Red	1,0,0	0°,1,1
Green	0,1,0	120°,1,1
Blue	0,0,1	240°,1,1
Yellow	1,1,0	60°,1,1
Cyan	0,1,1	180°,1,1
Magenta	1,0,1	300°,1,1
Sky Blue	.5, .5, 1	240°,.5,1

The CIE Chromaticity Diagram

- In 1931, an international commission met to develop a color model that could represent all visible colors *without negative amounts* of red, green or blue
- The primaries they derived are called X, Y and Z
- They are expressed using a set of *color matching formulas*
- The color space has a deformed cone-like shape
- The CIE chromaticity diagram is best used as a tool for studying other color spaces, like RGB and CMY

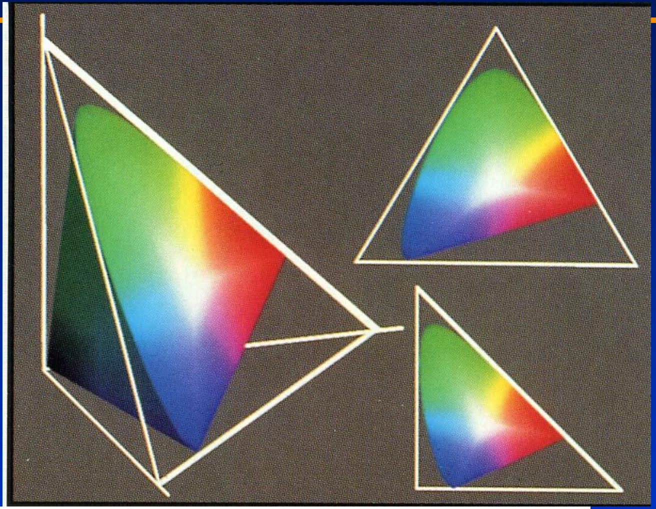

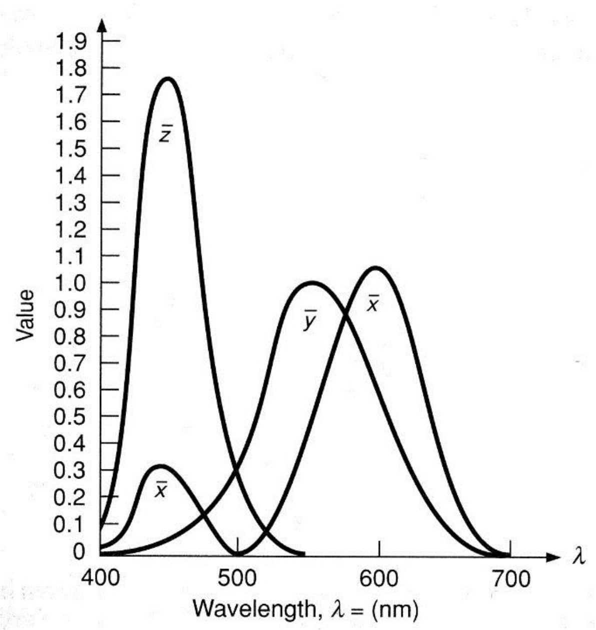

Uses of the CIE Chromaticity Diagram

- The CIE can tell us a lot about *color gamuts*, the range of colors that can be expressed with a particular color model

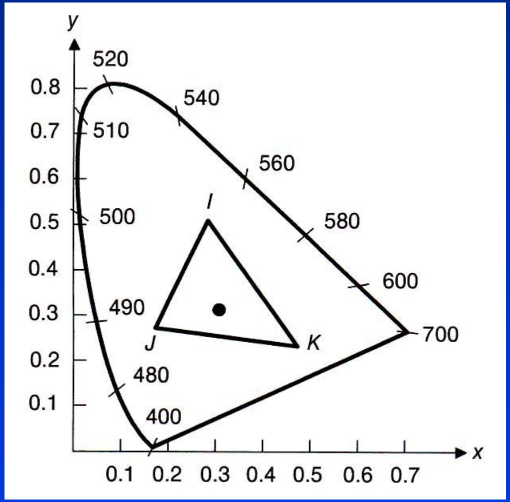

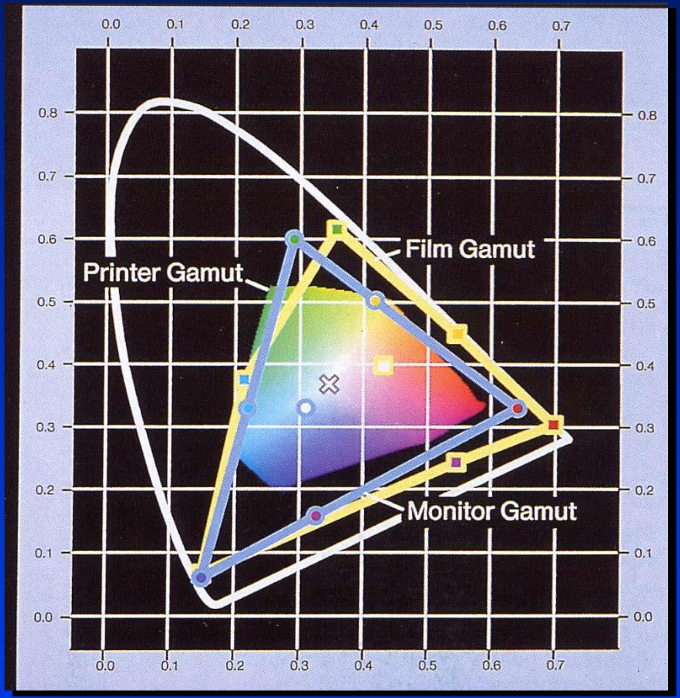

- This figure shows the color gamut for a hypothetical IJK color space
- Note that no single triangle can encompass the entire CIE diagram

- In other words, color monitors and printed materials cannot display the entire spectrum of possible colors

Human Eye Color Perception

- Color resolution - Human eye differentiates about 300 hues and 100-150 luminance variations.
- If red, green and blue cones are 60%, 30% and 10%, which colors can we perceive best?
- Best resolution is for green and red, less resolution for blue
- What does that mean?
- What does this mean for visualization?

Human Eye Color Response

- The time to response to a signal varies according to the color used
- Human eye responds to certain colors faster than others
- Color ranking (from best to worst):
  yellow > white > red > green > blue
- What colors should be used to highlight important features?
  – Important features should be visualized in light colors, such as yellow and white
- Background information is best visualized in dark colors, such as green and blue
- Let’s test your color response

Human Eye Color Perception

- Color channel properties:
  - luminance channel: detail, form, shading, stereo, motion
  - color: surfaces of things, labels, categories (about 10)
  - red, green, blue, yellow are special
- Chromatic channels have low resolution
  - luminance contrast needed to see detail (3:1 recommended, 10:1 for small text)

Some Natural philosophers suppose that these colors arise from accidental vapours diffused in the air, which communicate their own hues to the shadows; so that the colours of the shadows are occasioned by the reflection of any given sky colour: the above observations favour this opinion.

Text on an isoluminant background is hard to read

Information Coding with Color

- Color is very good for *classification* – separation of data into classes
- Color helps us determine type and character of data
- In practice, only about six categories can be distinguished using color alone
- Keep this in mind when developing your own visualization of datasets

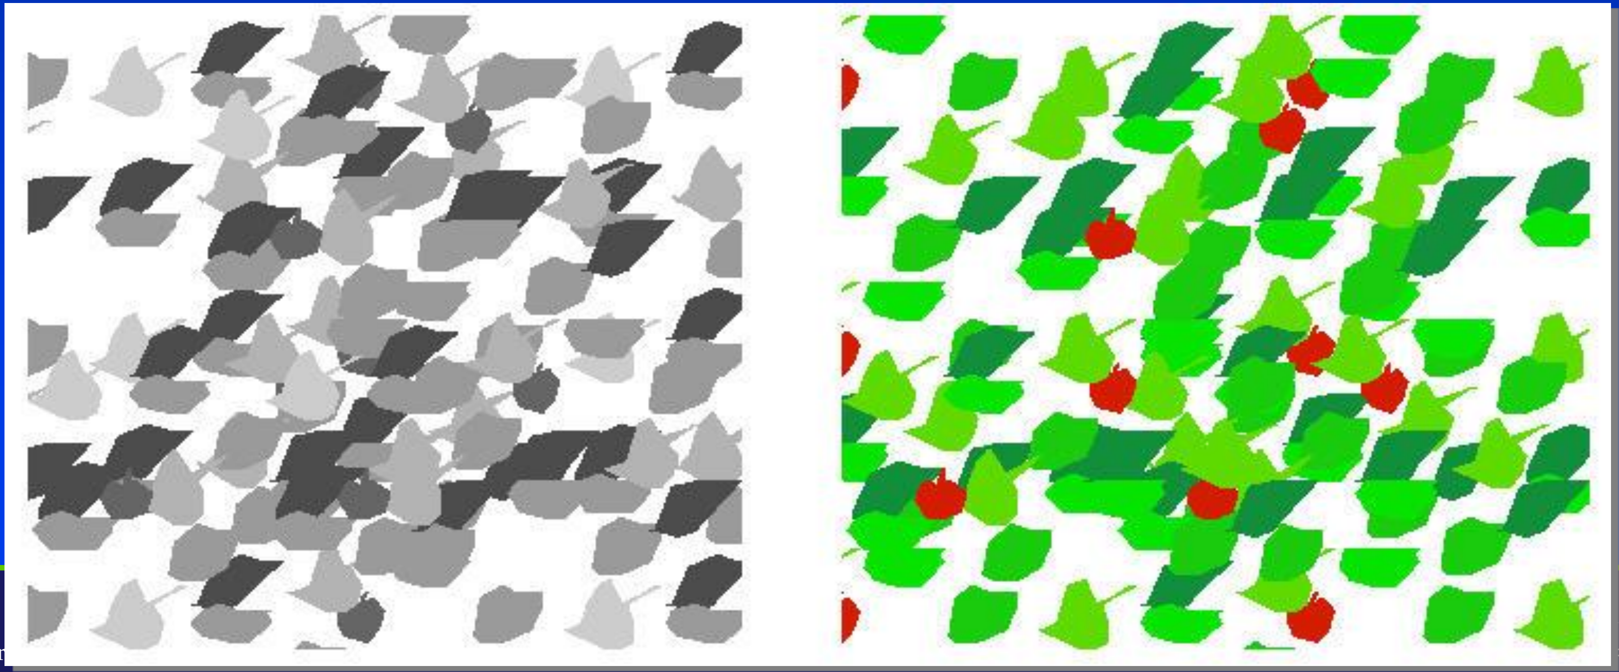

Information Coding with Color

- A suggested order for adding color to visualizations

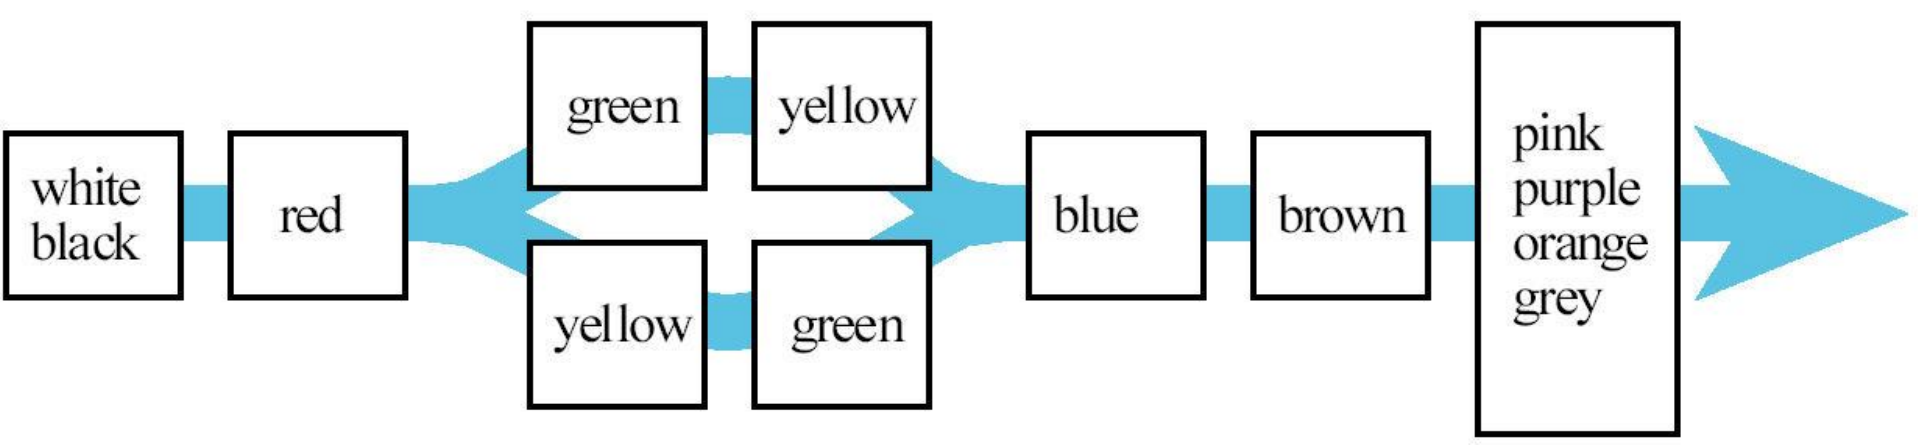

Information Coding with Color – Helpful Tips

- Color coding
  - large areas: low saturation
  - small areas: high saturation
  - maintain luminance contrast
  - break iso-luminances with borders (?)
- Example of these ideas in practice:

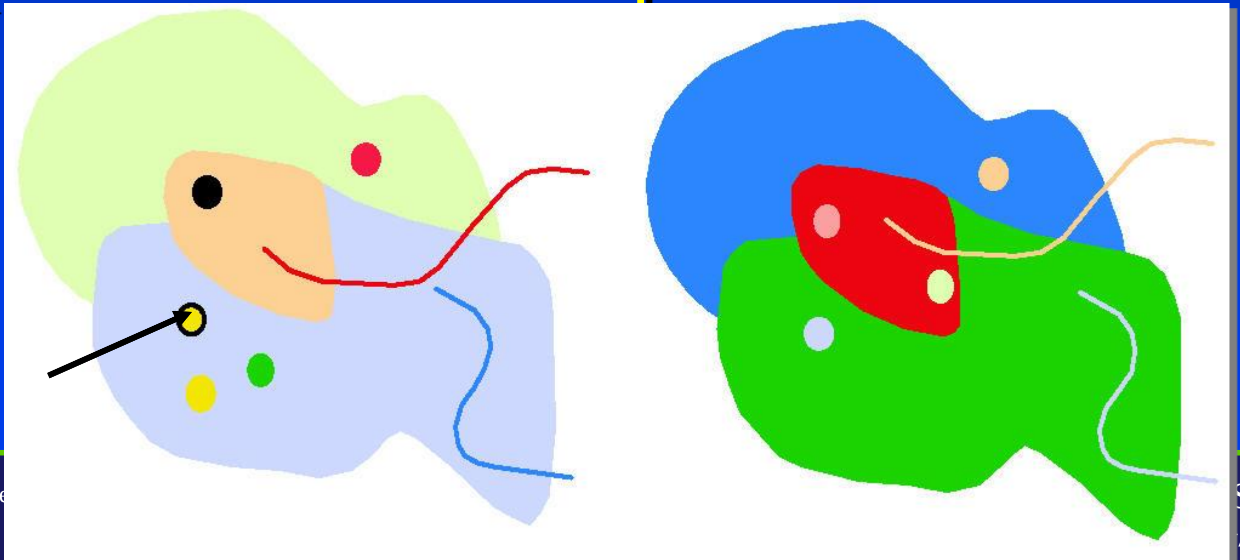

Human Eye Color Perception

- We are sensitive to small color differences
- Good at making side-by-side comparisons
- Not as good at identifying colors in isolation
- Hard: “Is that red, or orange-red, or red-orange, or maroon, or ….?”
- Easier: “Color A is redder than color B”

Color Response and Perception Examples

- Use bright colors to highlight important features
- Yellows, oranges, reds will be picked up first by the eye
- Make effective use of light/dark contrast:

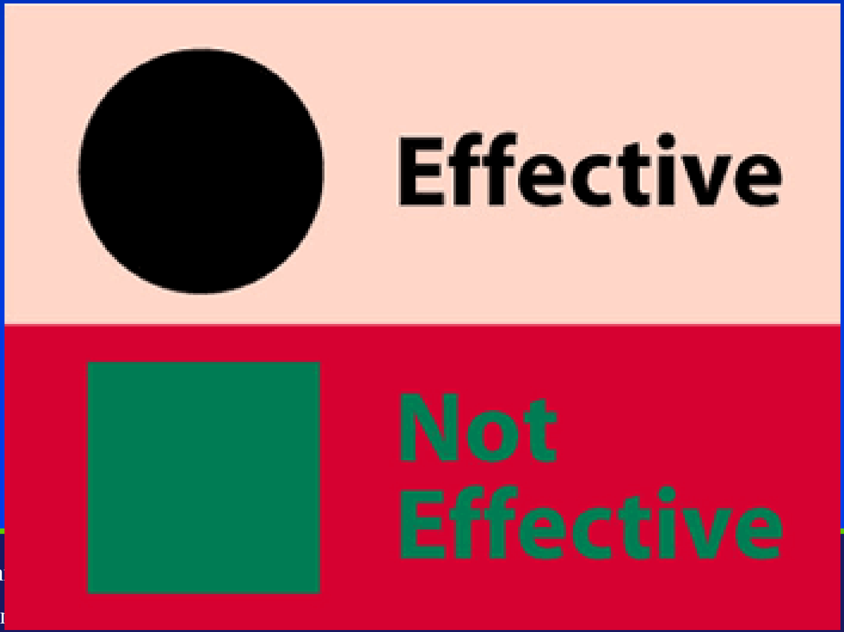

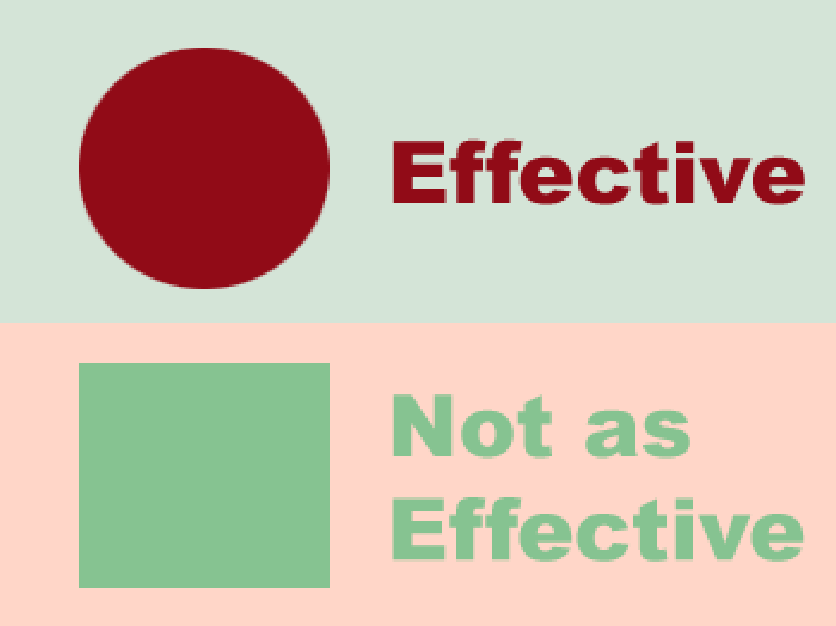

Color Response and Perception Examples

- The human eye is very good at making side-by-side comparisons
- Also make effective use of color contrast to highlight important and interesting characteristics of data

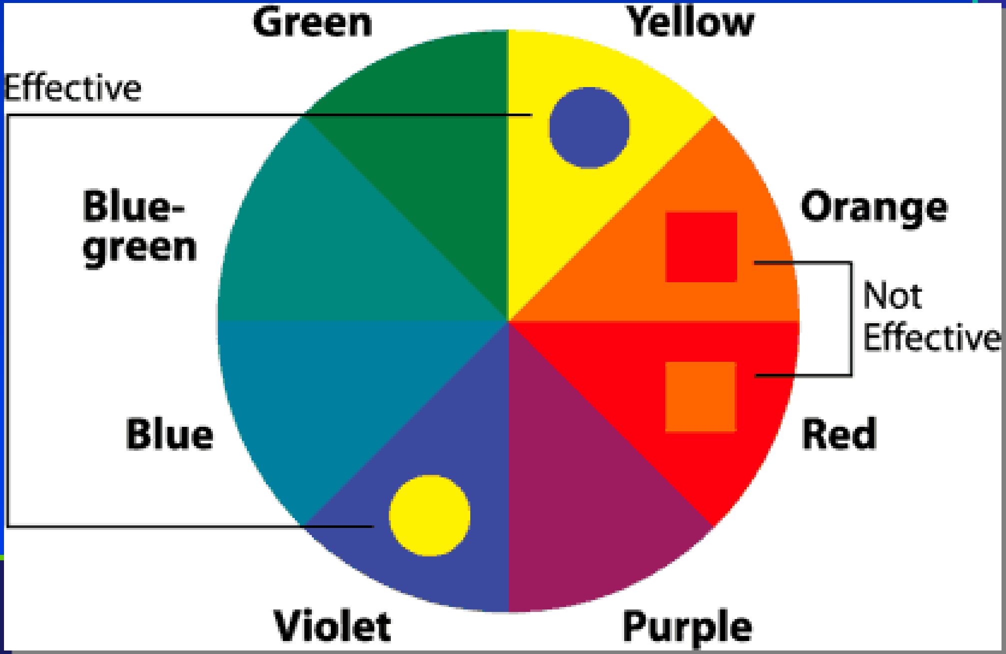

Color Perception

- Color blindness collapses the 3D color space into 2D
- 8% of males are color-blind
- Color resolution
  - color perception is relative
  - we are sensitive to small differences, meaning that we need millions of colors in practice
  - we are best able to distinguish colors when making side-by-side comparisons
  - however, we are not sensitive to absolute values, meaning that we can use only < 10 color values for coding information

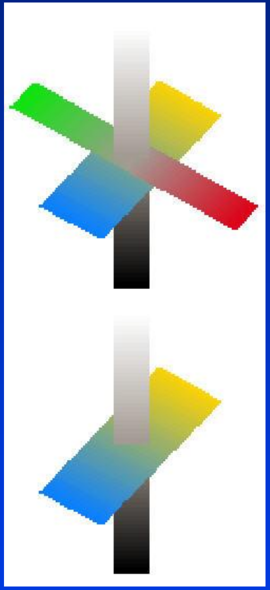

Pseudo-coloring (Colorization)

- Assign colors to gray levels by indexing the gray levels into a color map
- This indexing or translation takes place via a *transfer function*
- This is a simple yet powerful technique used in many types of visualization

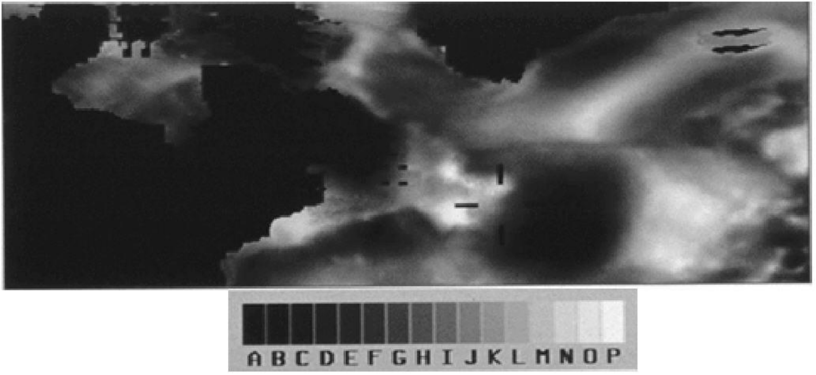

original graylevel map

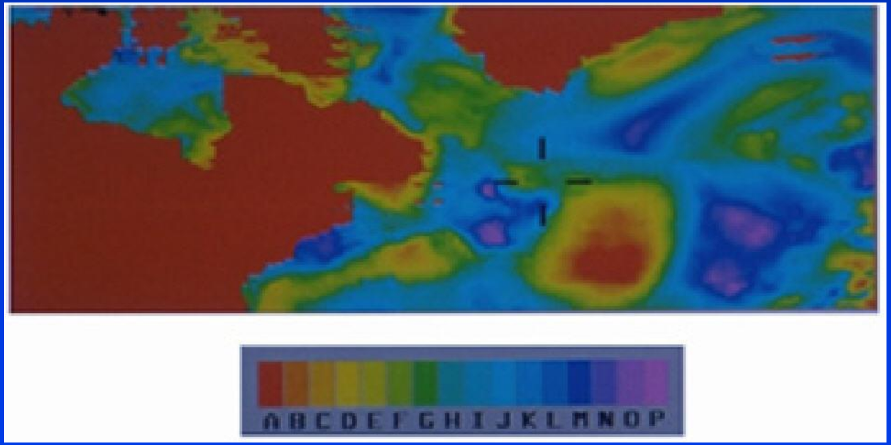

simple spectrum sequence
with iso-luminance

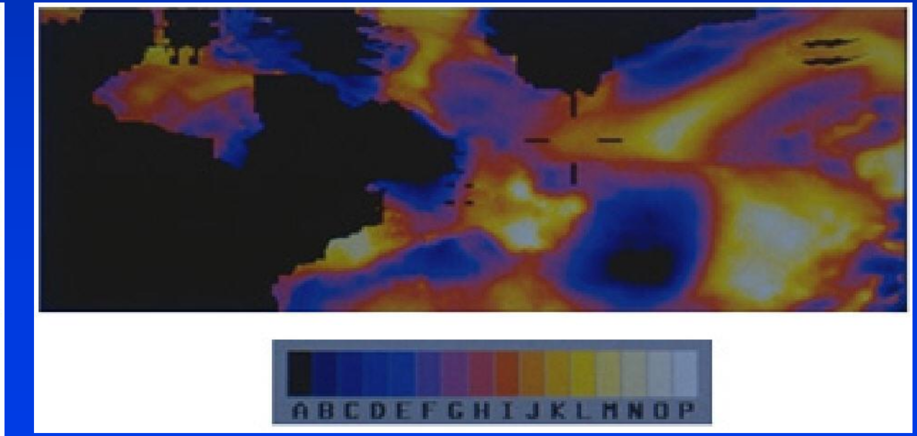

more effective: spiral sequence through
color space; luminance increases with hue

Color Perception - Summary

- Use luminance for detail, shape, and form
- Use color for coding - few colors
- Minimize contrast effects
- Strong colors for small areas - contrast in luminance with background
- Subtle colors can be used to segment large areas
- Figure at right shows how to use variation of luminance to indicate the direction of flow